

Probabilistic Structural Controllability in Causal Bayesian Networks

Ardavan S. Nobandegani[†] & Ioannis N. Psaromiligkos[†]

Abstract

Humans routinely confront the following key question which could be viewed as a probabilistic variant of the controllability problem: While faced with an uncertain environment governed by causal structures, how should they practice their autonomy by intervening on driver variables, in order to increase (or decrease) the probability of attaining their desired (or undesired) state for some target variable? In this paper, for the first time, the problem of probabilistic controllability in Causal Bayesian Networks (CBNs) is studied. More specifically, the aim of this paper is two-fold: (i) to introduce and formalize the problem of probabilistic structural controllability in CBNs, and (ii) to identify a sufficient set of driver variables for the purpose of probabilistic structural controllability of a generic CBN. We also elaborate on the nature of minimality the identified set of driver variables satisfies. In this context, the term “structural” signifies the condition wherein solely the structure of the CBN is known.

Introduction

The aptitude to perceive causation plays a central role in human cognition. Intervention, as a defining feature of humans to manifest their autonomy, is the sole means of actively (in contrast with the passive mode of being a mere observer) interacting with a world governed by causal structures. Among possible intentions behind exerting intervention, the notion of “control” is a notable one—that is, informally speaking, to manipulate some variables¹ (also called driver variables) of a system to, either directly or indirectly, “shape”, “guide”, or “control” variables of the system which are of interest.

In this work, the problem of Targeted Probabilistic Structural controllability (TPS-controllability) in the context of Causal Bayesian Networks (CBNs) is addressed and formalized. The term ‘structural’ signifies the condition wherein the agent is equipped merely with the causal structure of the domain under study. The term ‘targeted’, on the other hand, emphasizes that the agent is interested in controlling the behavior of a specific subset (or all) of variables in the domain

called target variables. Finally, the term ‘probabilistic’ highlights the probabilistic nature of the problem under study.

At a high level, we define the problem of *probabilistic controllability* in the context of CBNs as follows: *How an agent, provided with the knowledge of the set of intervenable variables, should devise her intervention, i.e., (Q.1) “which” variables to intervene on, and (Q.2) “how” to intervene on those, so as to “control” the behavior of a particular variable(s) of interest in the domain (represented by a CBN), that is, to maximize/minimize the probability of the occurrence of a state of interest for a set of target variables.*

The problem of probabilistic structural controllability in the context of CBNs is then accordingly defined as that of probabilistic controllability—as stated above—with one crucial additional constraint on the agent’s part: The agent is solely equipped with the knowledge of the underlying causal *structure* of the domain (i.e., the CBN’s topology) and is uninformed of the parametrization thereof. In this work, we aim at identifying the minimal set of intervenable variables sufficient for TPS-controllability of an *arbitrary* CBN. Particularly, we devise an algorithm, C^* , which identifies a sufficient set of intervenable variables for the purpose of TPS-controllability of a generic CBN. We also elaborate on the nature of minimality the identified set satisfies. Furthermore, to formalize the problem under study and to articulate the results, along the way: (i) we devise, building upon Pearl’s concept of *stochastic policy* (cf. (Pearl 2000), Sec. 4.2), a graphical representation of a generic intervention policy, and (ii) put forth a hierarchical construct for intervention policies wherein moving up in the hierarchy amounts to empowering the agent to exercise more advanced forms of intervention.

The question of interest to this work has significant ramifications for studies on strategic planning and policy making in domains enjoying causal structure. Equally importantly, the problem under study has notable connections to how humans, at the computational level of analysis (Marr 1982) and in line with the rational analysis approach (Anderson 1990), devise their intervention strategies efficiently to increase the odds of attaining their desired goals while faced with their uncertain environment.

Notation and Terminology

In this section, we present some preliminary notations and terminologies which will be adopted in this paper. Random

[†]The authors are with the Department of Electrical and Computer Engineering, McGill University.

¹The terms “node” and “variable” will be used interchangeably throughout.

quantities are denoted by bold-faced letters; their realizations are denoted by the same letter but non-bold. More specifically, Random Variables (RVs) are denoted by lower-case bold-faced letters, e.g., \mathbf{x} , and their realizations by non-bold lower-case letters, e.g., x . Likewise, sets of RVs are denoted by bold-faced calligraphic letters, e.g., \mathcal{X} , and their corresponding realizations by non-bold calligraphic letters, e.g., \mathcal{X} . $Val(\cdot)$ denotes the set of possible values a random quantity can take on. To simplify presentation, we incur the following abuse of notation: We denote the probability $\mathbb{P}(\mathbf{x} = x)$ by $\mathbb{P}(x)$ for some RV \mathbf{x} and its realization $x \in Val(\mathbf{x})$. For conditional probabilities, we will use the notation $\mathbb{P}(x|y)$ instead of $\mathbb{P}(\mathbf{x} = x|\mathbf{y} = y)$. Likewise, $\mathbb{P}(\mathcal{X}|\mathcal{Y}) := \mathbb{P}(\mathbf{X} = \mathcal{X}|\mathbf{Y} = \mathcal{Y})$ for $\mathcal{X} \in Val(\mathcal{X})$ and $\mathcal{Y} \in Val(\mathcal{Y})$. Random quantities are assumed to be discrete, throughout, unless stated otherwise.

Throughout the paper, the Directed Acyclic Graph (DAG) $G = (V, E)$ characterizes the non-intervened causal structure of the domain where V denotes the set of nodes/variables and E denotes the set of edges. We adopt Pearl’s notation $do(x) := do(\mathbf{x} = x)$ to denote an atomic intervention on \mathbf{x} so as to force it to take on value x . Also, $ip(\mathbf{x})$ denotes the intervention policy to be adopted for \mathbf{x} the meaning of which will be clarified in the subsequent section; informally intervention policy refers to how the agent decides to manipulate the intervened variable (cf. (Pearl 2000), Sec. 4.2). Intervention policy may or may not functionally depend on other variables of the domain. As we will see later, intervention policy in its most generic form is nothing but a Conditional Probability Distribution (CPD). $\delta(\cdot)$ denotes the Kronecker delta function. Also, Backward Chaining (BC) on a variable refers to the simple process of identifying its parents (i.e., immediate causes) and the parents of the parents and so forth until the boundaries of the CBN are reached.

Before proceeding further, let us formally define two key notions, namely, *subsumability* and *domination*.

Def. (Subsumability): DAG $G_1 = (V_1, E_1)$ *subsumes* DAG $G_2 = (V_2, E_2)$, denoted in short by $G_1 \supseteq G_2$, iff $V_1 = V_2$ and $E_2 \subseteq E_1$. We refer to the set $E_1 \setminus E_2$, as the *surplus* of G_1 with respect to G_2 .

Def. (Domination): DAG $G_1 = (V_1, E_1)$ *dominates* DAG $G_2 = (V_2, E_2)$, denoted by $G_1 \succcurlyeq G_2$, iff there *exists* a parametrization of G_1 which yields a result for the objective of interest that is no worse than the best achievable by *any* possible parametrizations of G_2 . For instance, if the objective of interest is to maximize the probability of some event of interest, say $\mathbf{r} = r$ for some $r \in Val(\mathbf{r})$, then we write $G_1 \succcurlyeq G_2$ iff there exists a parametrization of G_1 which yields some value for the probability of interest, $\mathbb{P}(r)$, which is greater than or equal to the best achievable by any possible parametrizations of G_2 .

Lemma 1. (Domination vs Subsumability): Let G_1, G_2 be DAGs. Then, $G_1 \supseteq G_2 \Rightarrow G_1 \succcurlyeq G_2$.

Proof: The proof is straightforward once we realize that one can very well take advantage of the extra edges of G_1 with respect to G_2 (i.e., the surplus of G_1 with respect to G_2) which give one more “degrees of freedom” to entertain and hence to achieve a result which is equally good or better

than the best achievable by all possible parametrizations of G_2 in terms of the objective of interest. ■

In subsequent sections where we introduce potential objectives of interest, the above statements become clearer.

Motivating Example

To develop some intuition as to the problem under study, a series of informative examples will be presented.




Figure 1: Motivating example.

Let us first consider the CBN depicted in Fig. 1(a). Variables $\mathbf{y}_1, \mathbf{y}_2$ are amenable to intervention (or, in short, *intervenable*). Let us assume that the objective is to make the occurrence of the event $\mathbf{o} = o \in Val(\mathbf{o})$ as likely as possible through intervening on any subset of variables $\{\mathbf{y}_1, \mathbf{y}_2\}$ (or to choose not to intervene at all corresponding to choosing the empty set). The key question is how, by mere investigation of the *structure* of the CBN depicted in Fig. 1(a), to decide: (i) on which (intervenable) variable(s) to intervene, and (ii) how the intervention should be exercised (a notion referred to as *Intervention Policy* (IP)). It is easy to come to the conclusion that, to make the occurrence of $\mathbf{o} = o$ as likely as possible, one needs to just intervene on \mathbf{y}_1 and force it into the state $\mathbf{y}_1 = y_1^*$ where y_1^* is the realization for \mathbf{y}_1 conditioned on which the probability of event $\mathbf{o} = o$ is maximum, i.e., $y_1^* = \arg \max_{y_1 \in Val(\mathbf{y}_1)} \mathbb{P}(o|y_1)$.

It is crucial to realize then that, due to the *structure* of the CBN depicted in Fig. 1(a), and *regardless* of its parametrization, it suffices for the agent to solely intervene on \mathbf{y}_1 for the purpose of TPS-controllability of $\mathbf{o} = o$ (the answer to (i)). Furthermore, since the agent is assumed to be equipped merely with the structure of the underlying CBN and not the parametrization thereof, based on the above argument on \mathbf{y}_1 ’s IP, the agent can just arrive at the conclusion that \mathbf{y}_1 ’s IP has the functional (or structural) form of² $\mathbb{P}(\mathbf{y}_1)$ —that is, merely the *non-parametric* form of the IP (the answer to (ii)). Altogether, a solution to the problem of TPS-controllability of $\mathbf{o} = o$ is $\{\mathbf{y}_1\}$ (which is a sufficient set of variables to be intervened) along with $\mathbb{P}(\mathbf{y}_1)$ which is the *functional* form of \mathbf{y}_1 ’s IP. Following the same line of reasoning for the CBNs depicted in Figs. 1(b-d), it is straightforward to argue that $\{\mathbf{y}_1\}$ is a sufficient set for TPS-

²In fact, the agent can reason out one step further and come to the conclusion that \mathbf{y}_1 ’s IP must have the functional form of $\mathbb{P}(\mathbf{y}_1) = \delta(\mathbf{y}_1 = y_1^*)$, however, the agent cannot identify/specify the value of y_1^* —due to the lack of knowledge about the parametrization of the CBN.

controllability of the target variable \mathbf{o} , and \mathbf{y}_1 's IP has the functional form of $\mathbb{P}(\mathbf{y}_1)$ akin to what we had for Fig. 1(a).




Figure 2: Motivating example.

Let us consider another example that highlights a key idea, namely, that we need to broaden our understanding of the notion of intervention (cf. (Pearl 2000), Sec. 4.2). Consider the CBN depicted in Fig. 2. This time, only \mathbf{y} is intervenable. Assume that (only for this particular example), all the variables are binary-valued; the prior probability on \mathbf{x} is $\mathbb{P}(\mathbf{x})$ (which is assumed to be non-degenerate), $\mathbf{y} := \neg \mathbf{x}$, and $\mathbf{o} := \mathbf{x} \oplus \mathbf{y}$ where \neg and \oplus denote the logical connectives *not* and *xor*, respectively. It is easy to verify that the event $\mathbf{o} = 1$ occurs with probability one, regardless of the choice of $\mathbb{P}(\mathbf{x})$. For the problem of TPS-controllability of $\mathbf{o} = 1$, the agent has to decide whether or not to intervene on \mathbf{y} . Imagine an intervention were to be exercised on \mathbf{y} . Whether the agent would set $\mathbf{y} = 0$, or $\mathbf{y} = 1$, the objective event of $\mathbf{o} = 1$ would become less likely to happen compared to that of the (non-intervened) original model. At first glance, the fact that exercising intervention makes the situation worse in terms of the objective of interest seems rather counter-intuitive. How could it be that having the freedom to manipulate variable \mathbf{y} (which even happens to be one of the parents of the objective node) whatever way we like does not allow you to outperform the (non-intervened) original model? The answer lies in developing a better understanding of the term “whatever way we like.” For that purpose, we need to broaden our conception of the notion of intervention and go beyond practicing merely a primitive atomic form of intervention denoted by $do(\mathbf{y} = y)$ in the literature. A more advanced form of intervention is to pick the state to which we want to force the intervened variable as a function of the states of some other variables of the domain. That is, IP may depend functionally on a collection of other variables in the domain. In this example, choosing \mathbf{y} 's IP, denoted by $ip(\mathbf{y})$, to functionally depend on (the state of) variable \mathbf{x} ensures that, the outcome achieved by exerting such intervention, is equally good or better (no worse than) than that of the (non-intervened) original model. In such a setting, we simply adopt the following terminology/notation: The intervention pair $(\mathbf{y}, ip(\mathbf{y}) := \mathbb{P}(\mathbf{y}|\mathbf{x}))$ (comprising, in order, of the set of intervened variables and their corresponding IPs) is equally good or better than both: (i) the intervention pair $(\mathbf{y}, ip(\mathbf{y}) := \mathbb{P}(\mathbf{y}) = \delta(\mathbf{y} = y))$ corresponding to the simplistic atomic form of intervention on \mathbf{y} discussed above which does not depend on the state of \mathbf{x} and, likewise, (ii) the intervention pair (\emptyset, \emptyset) corresponding to the original model (without intervention). A simple comparison between the CBN in Fig. 1(a) and the one in Fig. 2 reveals the following: Had there been left, upon intervening on \mathbf{y} , no paths through which \mathbf{x} 's causal effect on the objective vari-

able \mathbf{o} could be mediated, there would be no need for the intervention policy of \mathbf{y} to functionally depend on \mathbf{x} (akin to Fig. 1(a)).

Intervention Policy

The notion of IP delineates how an intervened variable should be “manipulated.” More specifically, IP indicates whether other variables play any role or not (and if so, how) in devising how the manipulation on a to-be-intervened variable is to be practiced. It is intuitive that the more variables we are allowed to functionally depend on while devising the IP of a to-be-intervened variable, the more “degrees of freedom” we have in controlling the behavior of the to-be-intervened variable. Next, we will formalize this intuition by, first, putting forward a simple graphical representation of intervention policy and, then, by using the idea of *subsumability* which draws on the fact that the original non-intervened model could be thought of as a special case of the intervened model. By allowing a larger number of variables for the IP of an intervened variable to *functionally* depend on, we show that IPs can be organized in a hierarchical construct wherein moving up in the hierarchy amounts to empowering the agent to exercise more sophisticated forms of intervention.

Hierarchical Construct

Before proceeding further let us make a definition: The scope of an IP is the set of variables, except the intervened variable itself, that the IP functionally depends on. For instance, for $ip(\mathbf{y}) := \mathbb{P}(\mathbf{y}|\mathbf{s})$, the scope is comprised of the variable \mathbf{s} . In short, the idea of organizing IPs into a hierarchical construct is inspired by the simple realization that, by delimiting the set of variables the agent is allowed to incorporate into the scope of the intervened variable's IP (functionally represented by a conditional probability distribution), we impose a constraint on the expressive power of the intervention policy.

The following notation henceforth will be employed to refer to different IP-classes:

- **IP class-0:** This class refers to the set of IPs where the scope of each is the empty set, i.e., a setting wherein the IP(s) of the intervened variable(s) is not allowed to incorporate any variables into its scope. That is, if variable \mathbf{x} happens to be decided to be intervened and the agent is only permitted to adopt *class-0* intervention policies, then, the agent is just allowed to place IP of the functional form $\mathbb{P}(\mathbf{x})$ on \mathbf{x} to exercise her intervention. It is crucial to note that, the agent is allowed to parameterize $\mathbb{P}(\mathbf{x})$ as she wishes, yet the functional form of the IP is constrained.

- **IP class-1:** This class refers to the set of IPs where the scope of each comprises the immediate causes of the corresponding intervened variable. That is, in this class, the IP of the intervened variable \mathbf{x} has the functional form $ip(\mathbf{x}) := \mathbb{P}(\mathbf{x}|\mathbf{par}(\mathbf{x}))$ where $\mathbf{par}(\mathbf{x})$ is the set of immediate causes for \mathbf{x} . That is, the scope of *class-1* IP $ip(\mathbf{x})$ is $\mathbf{par}(\mathbf{x})$.

- **IP class- j , $\forall j \geq 2$:** This class refers to the set of IPs where the scope of each is the ancestors of the corresponding intervened variable up to i^{th} level. For instance, for the

case of IP *class-2*, IP(s) of the intervened variable(s) is solely allowed to take into account the state of (i) the immediate causes, and (2) the immediate causes of the variables in (i), thereby, altogether functionally depending on all the ancestors up to the 2nd level.

- IP *class-∞*: This class refers to the set of IPs where the scope of each is *all* the ancestors of the corresponding intervened variable. Note that the complete set of ancestors of a variable \mathbf{x} can be found by instantiating the BC on \mathbf{x} .

Finally, it is crucial to notice the following. For an IP to be in a particular class amounts to imposing a constraint solely on the *functional* form of the IP; the agent is free to choose any parametrization for the IP as she may wish. Therefore, $ip(\mathbf{x}) \in \text{class-}i$ simply means that the *functional* form of $ip(\mathbf{x})$ is constrained in accord to the definition of IP *class-}i* given above, yet, it could be arbitrarily parameterized. Also, assuming that $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^m$, the notation $ip(\mathcal{X}) \in \text{class-}j$ will be adopted as a shorthand for the following: $ip(\mathbf{x}_i) \in \text{class-}j, \forall i = 1, \dots, m$.

Graphical Representation

In this section, we develop a way of visualizing IPs. The idea is simple. If the IP of a to-be-intervened variable \mathbf{x} happens to functionally depend on \mathbf{y} , then we show this by a directed dash-dotted arrow emanating from \mathbf{y} and pointing towards \mathbf{x} . To ensure that any practice of intervention is fully expressed by such edges we introduce the following convention: For DAG $G = (V, E)$, a clamped variable \mathfrak{C} is added³ to V . Then, intervening on a variable \mathbf{a} which has no parents and exerting $ip(\mathbf{a}) = \mathbb{P}(\mathbf{a})$ will be illustrated graphically by a dash-dotted edge emanating from \mathfrak{C} towards \mathbf{a} . In general, upon \mathbf{y} taking on the state y , the agent may decide to set the value of \mathbf{x} to a fixed value x (deterministic IP), or to set the value of \mathbf{x} probabilistically (stochastic IPs), i.e., \mathbf{x} takes on values from $Val(\mathbf{x})$ according to some conditional probability distribution $\mathbb{P}(\mathbf{x}|\mathbf{y})$. In both cases, $ip(\mathbf{x})$ is said to *functionally* depend on variable \mathbf{y} . Simply put, in devising the intervention policy of \mathbf{x} , namely, $ip(\mathbf{x})$, the state of \mathbf{y} is taken into account. The notion of probabilistic IP is discussed in ((Pearl 2000), pp. 113-114) under the title of *stochastic policy*.

Let us first give some definitions which will prove useful in the subsequent sections. For the given definitions, DAG $G = (V, E)$ represents the causal structure of the domain.

Def. (Intervention Pair): A set of intervened variables $\mathcal{K} \subseteq V$ along with their corresponding IPs comprise a pair, called an *intervention pair*, which is denoted by $(\mathcal{K}, ip(\mathcal{K}))$.

Def. (Intervention DAG (i-DAG)): Every intervention pair $(\mathcal{K}, ip(\mathcal{K}))$ for $\mathcal{K} \subseteq V$ uniquely specifies a DAG (as described above) which we refer to as the *i-DAG* associated to that intervention pair. The *i-DAG* associated to $(\mathcal{K}, ip(\mathcal{K}))$ is denoted by $(\mathcal{K}, ip(\mathcal{K}))_G$.

Def. (i-Subsumability): For $\mathcal{X}, \mathcal{Y} \subseteq V$, *i-DAG* $(\mathcal{X}, ip(\mathcal{X}))_G$ *i-subsumes* *i-DAG* $(\mathcal{Y}, ip(\mathcal{Y}))_G$, denoted in short by $(\mathcal{X}, ip(\mathcal{X}))_G \supseteq_i (\mathcal{Y}, ip(\mathcal{Y}))_G$, iff (i)

³It is implicitly assumed throughout this paper that the variable \mathfrak{C} has been added to DAG G *a priori*. Also, we will not depict \mathfrak{C} in the figures unless needed.




Figure 3: Sample case. (a): Original CBN. Variable \mathbf{y} is to be intervened according to $ip(\mathbf{y}) := \mathbb{P}(\mathbf{y}|\mathbf{u})$. (b): The graphical representation of intervening on \mathbf{y} with $ip(\mathbf{y}) := \mathbb{P}(\mathbf{y}|\mathbf{u})$. The figure simply illustrates the fact that the state of \mathbf{y} gets decided (potentially probabilistically) according to the state of \mathbf{u} . Notice that according to (Pearl 2000), upon intervening on \mathbf{y} , all the incoming edges into \mathbf{y} should be removed.

($\mathcal{X}, ip(\mathcal{X}))_G \supseteq (\mathcal{Y}, ip(\mathcal{Y}))_G$, (ii) the set of the dash-dotted edges in $(\mathcal{Y}, ip(\mathcal{Y}))_G$ is a subset of the set of the dash-dotted edges in $(\mathcal{X}, ip(\mathcal{X}))_G$, and (iii) the surplus of $(\mathcal{X}, ip(\mathcal{X}))_G$ with respect to $(\mathcal{Y}, ip(\mathcal{Y}))_G$ is solely comprised of dash-dotted edges.

Def. (i-Domination): For $\mathcal{X}, \mathcal{Y} \subseteq V$, *i-DAG* $(\mathcal{X}, ip(\mathcal{X}))_G$ *i-dominates* *i-DAG* $(\mathcal{Y}, ip(\mathcal{Y}))_G$, denoted by $(\mathcal{X}, ip(\mathcal{X}))_G \succ_i (\mathcal{Y}, ip(\mathcal{Y}))_G$ for short, iff there exist a parameterization for the dash-dotted edges (see Fig. 3) in $(\mathcal{X}, ip(\mathcal{X}))_G$ which yields a result for the objective of interest that is no worse than the best achievable by any possible parameterizations of the dash-dotted edges in *i-DAG* $(\mathcal{Y}, ip(\mathcal{Y}))_G$.

Lemma 2. (i-Domination vs i-Subsumability): Let DAG $G = (V, E)$ characterize the causal structure of the domain. Let $(\mathcal{X}, ip(\mathcal{X}))_G$ and $(\mathcal{Y}, ip(\mathcal{Y}))_G$ be two *i-DAGs* for some $\mathcal{X}, \mathcal{Y} \subseteq V$. Then, the following holds:

$$(\mathcal{X}, ip(\mathcal{X}))_G \supseteq_i (\mathcal{Y}, ip(\mathcal{Y}))_G \Rightarrow (\mathcal{X}, ip(\mathcal{X}))_G \succ_i (\mathcal{Y}, ip(\mathcal{Y}))_G.$$

Proof: The rationale is similar to the one presented for Lemma 1. ■

It immediately follows (due to Lemma 2) that, for $G = (V, E)$ and $\forall \mathcal{K} \subseteq V$, the following holds true: $\forall j \geq m$, $(\mathcal{K}, ip(\mathcal{K})) \in \text{class-}j)_G \succ_i (\mathcal{K}, ip(\mathcal{K})) \in \text{class-}m)_G$.

TPS-Controllability of CBNs: Formalization

Let DAG $G = (V, E)$ characterize the causal structure of the domain. Let $V = V_i \cup \bar{V}_i$ where V_i denotes the set of nodes/variables amenable to intervention (or, in short, *intervenable*), and \bar{V}_i be the complement of V_i , i.e., $\bar{V}_i = V \setminus V_i$. The probability of interest, in its generic form takes the form⁴ $\mathbb{P}(\mathcal{O} = \mathcal{O} | do[\mathcal{X}; ip(\mathcal{X}) = ip(\mathcal{X})])$, or in short $\mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})])$, where \mathcal{O} denotes the set of target variables, \mathcal{O} denotes the realization of interest, \mathcal{X} denotes the set of intervened variables, $ip(\mathcal{X})$ denotes the *functional* form (i.e., non-parametric representation) of the to-be-adopted intervention policy, and $ip(\mathcal{X})$ denotes a specific parametrization⁵ of $ip(\mathcal{X})$. Also, $do[\mathcal{X}; ip(\mathcal{X})]$ de-

⁴The connection to Pearl's notation for *do*-calculus is as follows (cf. (Pearl 2000), p. 114): $\mathbb{P}(y) |_{\mathbb{P}^*(x|z)} = \mathbb{P}(y | do[\mathbf{x}; \mathbb{P}^*(x|z)])$.

⁵Which is equivalent to a specific parametrization of the dash-dotted edges representing the intervention policy exercised on vari-

notes the setting wherein variables \mathcal{X} are intervened according to $ip(\mathcal{X})$ (functional form) and, likewise, $do[\mathcal{X}; ip(\mathcal{X})]$ denotes the setting wherein variables \mathcal{X} are intervened according to $ip(\mathcal{X}) = ip(\mathcal{X})$. One can write, $\mathcal{O} = \mathcal{O}_i \cup \bar{\mathcal{O}}_i$ where $\mathcal{O}_i \subseteq V_i$ and $\bar{\mathcal{O}}_i \subseteq \bar{V}_i$. Objectives of interest could have any of the following forms:

$$\min_{\mathcal{X} \subseteq V_i} \left(\min_{ip(\mathcal{X}) \in \text{class-}\infty} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right), \quad (1)$$

$$\max_{\mathcal{X} \subseteq V_i} \left(\max_{ip(\mathcal{X}) \in \text{class-}\infty} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right), \quad (2)$$

and,

$$\min_{\mathcal{X} \subseteq V_i} \left(\max_{ip(\mathcal{X}) \in \text{class-}\infty} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right), \quad (3)$$

$$\max_{\mathcal{X} \subseteq V_i} \left(\min_{ip(\mathcal{X}) \in \text{class-}\infty} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right). \quad (4)$$

In the sequel, we focus on objectives (1) and (2). The discussion on objectives (3) and (4) is deferred to the end of the paper. Therefore, whenever we use the statement “objective of interest” we are specifically referring to both of objectives (1) and (2) unless stated otherwise.

Next, we devise an algorithm, \mathcal{C}^* , for the problem of TPS-controllability of CBNs. \mathcal{C}^* outputs a set of intervenable variables, \mathcal{X}^* , which is “optimal” with respect to objectives (1) and (2). In other words, \mathcal{X}^* is a sufficient choice of variables to intervene on (according to IP *class- ∞*) to satisfy objectives (1) and (2). To put it formally, in the next section we show that for \mathcal{X}^* the following holds: $\forall \mathcal{Y} \subseteq V_i$,

$$(\mathcal{X}^*, ip(\mathcal{X}^*) \in \text{class-}\infty)_G \succsim_i (\mathcal{Y}, ip(\mathcal{Y}) \in \text{class-}\infty)_G.$$

TPS-Controllability of CBNs: Algorithm \mathcal{C}^*

Let us explain simply how \mathcal{C}^* works. BC has to be initiated on nodes in \mathcal{O} . Upon reaching any node in V_i , the BC execution path terminates at that node. This procedure continues until, for all of the BC execution paths, either: (i) The BC execution path gets terminated at some node belonging to V_i , or (ii) a node with no parents is reached. The set of intervenable variables at which BC terminates constitute \mathcal{C}^* 's output denoted by \mathcal{X}^* . Fig. 4 depicts a sample execution of \mathcal{C}^* .

Algorithm \mathcal{C}^* : Justification

On the Sufficiency of \mathcal{X}^* : Let us now present the sufficiency proof for \mathcal{X}^* where by sufficiency we mean that: intervening (according to IP *class- ∞*) on any variables in addition to \mathcal{X}^* does not yield any improvement upon what is achievable through merely intervening (according to IP *class- ∞*) on \mathcal{X}^* .

Notice that we can write $V_i = V_i^{BC} \cup \bar{V}_i^{BC}$ where V_i^{BC} is the set of intervenable variables which belong to the sub-graph generated by executing BC on the target nodes \mathcal{O} . It is obvious that intervening on any variable in V_i^{BC} is pointless

ables \mathcal{X} in the corresponding *i*-DAG (see Fig. 3).




Figure 4: Sample Case: Variable \mathbf{o} (depicted in red) is the target variable. The intervenable variables (i.e., members of V_i) are circled. The BC execution paths are colored in blue and illustrated by dash-dotted lines. Upon initiating the BC at the target variable \mathbf{o} , we arrive at the \mathbf{t}_1 (depicted in purple) located at the junction. Next, we arrive at \mathbf{t}_2 and \mathbf{t}_3 . Since $\mathbf{t}_3 \in V_i$, the BC terminates at \mathbf{t}_3 . On the other hand, since $\mathbf{t}_2 \notin V_i$, the BC continues. Having performed the BC on \mathbf{t}_2 , we arrive at \mathbf{t}_4 and \mathbf{t}_5 . Since $\mathbf{t}_4 \in V_i$, the BC terminates on \mathbf{t}_4 . At the end, since \mathbf{t}_5 (depicted in grey) has no parents (immediate causes), the BC terminates at \mathbf{t}_5 as well. Therefore, by mere investigation of the structure, \mathcal{C}^* outputs the set $\mathcal{X}^* = \{\mathbf{t}_3, \mathbf{t}_4\}$ as a solution to the objectives (2) and (1) for this particular setting (i.e., the given CBN and the corresponding V_i and \mathbf{o} as the target variable).

due to the following argument⁶: (\star) variables in \bar{V}_i^{BC} have no direct or indirect causal effect on any of the target nodes. Now, what is left to be shown is why, among all V_i^{BC} , it suffices to intervene on \mathcal{X}^* (according to IP *class- ∞*) and, why, intervening on any additional variables does not improve upon what is achievable through intervening merely on \mathcal{X}^* . More formally, the question is why the following holds: $\forall \mathcal{Y} \subseteq V_i^{BC}$

$$(\mathcal{X}^*, ip(\mathcal{X}^*) \in \text{class-}\infty)_G \succsim_i (\mathcal{Y}, ip(\mathcal{Y}) \in \text{class-}\infty)_G.$$

Notice that, the extension to the case of $\forall \mathcal{Y} \subseteq V_i$ immediately follows from argument (\star) . The realization of the fact that variable $\mathbf{y} \in \bar{V}_i^{BC}$ was not selected (for intervention) by \mathcal{C}^* implies that \mathbf{y} 's causal effect on the target variables which are descendant⁷ of \mathbf{y} must have been mediated through some of the selected nodes by \mathcal{C}^* say $\mathcal{Y}^\dagger \subseteq \mathcal{X}^*$ (otherwise, \mathbf{y} would have been selected). The claim as to the redundancy of further intervening on \mathbf{y} in addition to exerting intervention (according to IP *class- ∞*) on \mathcal{Y}^\dagger is as follows. Let $G = (V, E)$ be the DAG associated to the (non-intervened) underlying causal structure of the domain. First, notice that the effect of stochastic policies can be expressed in terms of atomic interventions in G as explained in ((Pearl 2000), pp. 113-114) and ((Pearl 1995), p. 684). Due to Rule 2 of Pearl's *do*-calculus, intervention on \mathbf{y} can be exchanged with \mathbf{y} being merely passively observed. Subsequently, due to Rule 1 of Pearl's *do*-calculus, \mathbf{y} 's observation can be dismissed. Notice that intervening on \mathcal{Y}^\dagger according to IP *class-*

⁶This statement immediately follows from Rule 3 of Pearl's *do*-calculus (cf. (Pearl 2000), p. 95).

⁷Intervening on \mathbf{y} , obviously, could merely influence \mathbf{y} 's descendant.

∞ (which amounts to incorporating all the ancestors of \mathbf{y}^\dagger into their IPs) renders \mathbf{y} d -separated from the target variables which are descendant of⁸ \mathbf{y} , hence the applicability of Rule 2 and Rule 1. This concludes the proof.

On the Minimality of \mathcal{X}^*

Let us present two definitions and a proposition which bears on the minimality of \mathcal{X}^* .

Def. (Locally Structurally Minimal (LSM)): \mathcal{C}^* 's output, \mathcal{X}^* , is LSM with respect to DAG G iff there exists a parametrization of G such that *no* proper subset of \mathcal{X}^* , namely, \mathcal{X}^{**} , exists for which the following holds:

$$(\mathcal{X}^{**}, ip(\mathcal{X}^{**}) \in class-\infty)_G \succ_i (\mathcal{X}^*, ip(\mathcal{X}^*) \in class-\infty)_G.$$

Def. (Uniformly Structurally Minimal (USM)): \mathcal{C}^* 's output, \mathcal{X}^* , is USM iff \mathcal{X}^* is LSM with respect to any DAG G .

Proposition 1. \mathcal{X}^* is USM with respect to objective (1).

Proof: The proof is constructive. The objective of interest is maximax given in (1). Let us assume, without loss of generality, that all the RVs are binary-valued and the desired state is for all the target variables to take on value one. Our goal is to parameterize an *arbitrary* G in such a way that: (i) the desired state happens with probability one if variables \mathcal{X}^* are all set to one through exerting atomic interventions, and (ii) the desired state happens with probability zero otherwise. Start at \mathcal{X}^* . Parameterize the CPD of each $\mathbf{x}^* \in \mathcal{X}^*$ such that it always takes on the value zero. Moving along the BC execution paths terminated at \mathcal{X}^* , proceed towards the target variables which are descendants⁹ of \mathcal{X}^* . Along the way, parameterize the CPD associated to any variable \mathbf{k} such that, conditioned on \mathbf{k} 's parents which are descendants of \mathcal{X}^* (denoted by $par_{\mathcal{X}^*}(\mathbf{k})$), \mathbf{k} takes on value one iff all $par_{\mathcal{X}^*}(\mathbf{k})$ take value one¹⁰. In other words, intermediate variables like \mathbf{k} work as an *and* logical gate. Proceed in the aforementioned manner until all the target variables (which are descendants of \mathcal{X}^*) are reached. It is easy to verify that indeed the desired state happens with probability one iff all the variables \mathcal{X}^* are set to one and, furthermore, intervening on any proper subset of \mathcal{X}^* in any way does not yield such an outcome. This concludes the proof. ■

On Minimax/Maximin Objectives

In this section we claim that, subject to the constraint that the IP's of the to-be-intervened variables has to belong to IP $class-j$, the solution to both minimax and maximin problem is the empty set, for all $j \geq 1$. Let us present a lemma using which it is easy to justify the claim made above.

Lemma 3. Let DAG $G = (V, E)$ represent the causal structure of the domain. $\forall j \geq 1$ and $\forall \mathcal{X} \subseteq V_i$, the following

⁸Incorporation of all the antecedents of \mathbf{y}^\dagger into their IPs guarantees that all the back-door paths from \mathbf{y} to the target variables which are descendant of \mathbf{y} are blocked (cf. (Pearl 2000), Sec. 3.3.1).

⁹For any target variable \mathbf{q} which is not a descendant of \mathcal{X}^* , parameterize $\mathbb{P}(\mathbf{q}|par(\mathbf{q}))$ such that \mathbf{q} takes the value 1 with probability one.

¹⁰In other words, the CPD of $\mathbb{P}(\mathbf{k}|par(\mathbf{k}))$ is parameterized in such a manner that the parents of \mathbf{k} which are *not* descendants of \mathcal{X}^* are rendered ineffective.

inequalities hold:

$$\begin{aligned} \min_{ip(\mathcal{X}) \in class-j} \mathbb{P}(\mathcal{O}|do[\mathcal{X}; ip(\mathcal{X})]) &\leq \mathbb{P}(\mathcal{O}), \\ \mathbb{P}(\mathcal{O}) &\leq \max_{ip(\mathcal{X}) \in class-j} \mathbb{P}(\mathcal{O}|do[\mathcal{X}; ip(\mathcal{X})]). \end{aligned}$$

The proof for Lemma 3 is straightforward due to the simple realization that, $\forall j \geq 1$, $(\mathcal{X}, ip(\mathcal{X}) \in class-j)_G$ i -subsumes the original CBN, thus $(\mathcal{X}, ip(\mathcal{X}) \in class-j)_G \succ_i (\emptyset, \emptyset)_G$, $\forall j \geq 1$. Using Lemma 3, it is easy to justify the claim we made earlier as to the solution to the minimax and maximin problems; the solution to both is the empty set. For details, the reader is referred to Appendix.

Related Work

In this section, we give an overview of the ideas explored in the literature which are, in spirit, related to the problem under study in this work. The idea of Structure Control Theory (SCT) proposed by Lin (Lin 1974) in the context of Linear Time-Invariant (LTI) systems governed by first-order differential equations (a.k.a. state equations) perhaps comes closest to our problem. In such domains, all variables are deterministic and the states of variables change in time according to the dynamics represented by state equations.

Authors in (Liu, Slotine, and Barabási 2011), drawing on the idea of SCT proposed by Lin (Lin 1974), aim at identifying the minimal set of variables which are sufficient for the purpose of structural controllability of a generic large-scale LTI system. In a subsequent work, authors in (Gao et al. 2014), relax the objective of structural controllability of the system in whole, to merely that of a particular set of desired variables called target variables. This line of thought has been motivated due to the understanding that in large-scale systems, it may neither be attainable nor required to control the full system but, rather, to merely control a subset of the variables of the system (analogous to target variables in our problem) which are deemed pivotal for the realization of the task at hand. In this light, (Gao et al. 2014) is concerned with the very same question underlying our work, yet, perusing it in radically different settings. In (Liu, Slotine, and Barabási 2011; Gao et al. 2014) both variables and their inter-connections are deterministic in nature whereas, in our case, both have probabilistic natures; a point of departure which leads to a substantially different line of work—both semantically and syntactically.

Conclusion

In this paper, for the first time, the problem of TPS-controllability in the context of CBNs was investigated and formalized. Algorithm \mathcal{C}^* was devised to identify a sufficient set of intervenable variables for the purpose of TPS-controllability of a generic CBN; the minimality of \mathcal{C}^* 's output was also characterized. The provided results can have significant ramifications for studies on strategic planning and policy making allowing one to efficiently practice her interventions in order to maximize (minimize) the odds of the desired (undesired) outcomes.

References

- Anderson, J. R. 1990. *The adaptive character of thought*. Lawrence Earlbaum Associates, Inc.
- Gao, J.; Liu, Y.-Y.; D’Souza, R. M.; and Barabási, A.-L. 2014. Target control of complex networks. *Nature communications*, 5:5415.
- Glymour, C. 2003. Learning, prediction and causal bayes nets. *Trends in cognitive sciences*, 7(1):43–48.
- Krynski, T. R., and Tenenbaum, J. B. 2007. The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*, 136(3):430.
- Lin, C. T. 1974. Structural controllability. *IEEE Transactions on Automatic Control*, 19(3):201–208.
- Liu, Y.-Y.; Slotine, J.-J.; and Barabási, A.-L. 2011. Controllability of complex networks. *Nature*, 473(7346):167–173.
- Marr, D. 1982. Vision: A computational approach.
- Pearl, J. 1995. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 1st edition.

Appendix

On Minimax/Maximin Objectives: Details Using Lemma 3 in the paper, the next inequalities follows. $\forall j \geq 1$,

$$\max_{\mathcal{X} \subseteq V_i} \left(\min_{ip(\mathcal{X}) \in \text{class-}j} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right) \leq \mathbb{P}(\mathcal{O}),$$

$$\mathbb{P}(\mathcal{O}) \leq \min_{\mathcal{X} \subseteq V_i} \left(\max_{ip(\mathcal{X}) \in \text{class-}j} \mathbb{P}(\mathcal{O} | do[\mathcal{X}; ip(\mathcal{X})]) \right).$$

One can readily conclude from the above inequalities the claim made in the paper as to the solution to the minimax and maximin problems; the solution to both is the empty set.

Minimality of \mathcal{X}^* for Minim Objective We present the result in a form of a proposition.

Proposition A.1: *Let DAG $G = (V, E)$ characterize the causal structure of the domain and $\mathcal{O}_i \neq \emptyset$. Any set $\{\mathbf{o}^*\}$ where $\mathbf{o}^* \in \mathcal{O}_i$ is a solution to the minimin objective given in (2).*

Proof: Setting $\mathbf{o}^* \in \mathcal{O}_i$ to any value other than the desired one simply renders the desired event $\mathcal{O} = \emptyset$ impossible which subsequently yields value zero for the minimin objective. This concludes the proof. ■

On the Connection to Psychology

Algorithm \mathcal{C}^* captures one very common-sensical strategy an individual (referred to as the agent throughout in the paper) follows while faced with the task under study: “The closer a to-be-intervened node is to the target node(s), the better.” That is, there exists this tendency to intervene on the nodes which are closest¹¹ to the target node(s). We refer to this intuitive tendency as the *proximity principle*. In psychology literature (Glymour 2003; Krynski and Tenenbaum 2007), CBN is adopted as a normative model, at the computational level of analysis (Marr 1982), to represent and reason about causal domains. Algorithm \mathcal{C}^* , therefore, could be viewed as a potential proposal at Marr’s algorithmic level of analysis (Marr 1982), for the problem of how subjects devise their intervention strategies to “control” the state of a target node. The plausibility of \mathcal{C}^* as a potential algorithmic-level candidate is elevated due to \mathcal{C}^* ’s capturing nicely, by its mere design (i.e., starting at target nodes and moving backwards toward intervenable nodes), the common-sensical proximity principle.

¹¹The extreme case being to intervene right on the target node(s) and set it to the desired value.